

UNCG COMPUTER SCIENCE COLLOQUIUM SERIES



Reinforcement Learning on Batch Job Scheduling

GUEST SPEAKER: DR. DONG DAI
WED, OCTOBER 27TH, 10:00-11:00 AM

ONLINE:
<https://go.uncg.edu/cs-colloq>

Abstract

High-performance computing (HPC) platforms leverage the centralized job schedulers to prioritize and schedule the waiting batch jobs. As HPC has multiplied in scale and complexity in the last decay, traditional heuristic methods face challenges in handling the highly dynamic scheduling environment. As a result, the need for more efficient job scheduling has also increased significantly. In this talk, I will describe our responses to such a challenge. Specifically, we propose leveraging the latest machine learning techniques, particularly deep reinforcement learning (DRL) techniques, to design new and optimize existing HPC job scheduling. We mainly explore two directions: 1) the feasibility and challenges of building an automated HPC job scheduler entirely using reinforcement learning; 2) the feasibility of using reinforcement learning to optimize existing HPC job schedulers. Our results show in both cases, RL can successfully learn the complex patterns of the scheduling problem and deliver better performance than state-of-the-art HPC schedulers.

Bio

Dong Dai is an Assistant Professor in the Computer Science Department at UNC Charlotte. He received his Ph.D in computer science from University of Science and Technology of China. Before joining UNC Charlotte, he was a research assistant professor at Texas Tech University. His primary research interests are in the general area of data-intensive high-performance computing systems. He works on intelligent infrastructure and algorithms to optimize the performance and strengthen the reliability of high-performance computing (HPC) systems.